

Milking the spherical cow — on aspherical dynamics in spherical coordinates

Andrew Pontzen^{1,2,3,4}, Justin I. Read⁵, Romain Teyssier⁶, Fabio Governato⁷,
Alessia Gualandris⁵, Nina Roth¹, Julien Devriendt²

¹ *Department of Physics and Astronomy, University College London, London WC1E 6BT*

² *Oxford Astrophysics, Denys Wilkinson Building, Keble Road, Oxford, OX1 3RH*

³ *Balliol College, Broad Street, Oxford, OX1 3BJ*

⁴ *Email: a.pontzen@ucl.ac.uk*

⁵ *Department of Physics, University of Surrey, Guildford, GU2 7XH, Surrey, UK*

⁶ *Institute for Theoretical Physics, University of Zurich, CH-8057 Zürich, Switzerland*

⁷ *Astronomy Department, University of Washington, Seattle, WA 98195, US*

Received —; published —.

ABSTRACT

Galaxies and the dark matter halos that host them are not spherically symmetric, yet spherical symmetry is a helpful simplifying approximation for idealised calculations and analysis of observational data. The assumption leads to an exact conservation of angular momentum for every particle, making the dynamics unrealistic. But how much does that inaccuracy matter in practice for analyses of stellar distribution functions, collisionless relaxation, or dark matter core-creation?

We provide a general answer to this question for a wide class of aspherical systems; specifically, we consider distribution functions that are “maximally stable”, *i.e.* that do not evolve at first order when external potentials (which arise from baryons, large scale tidal fields or infalling substructure) are applied. We show that a spherically-symmetric analysis of such systems gives rise to the false conclusion that the density of particles in phase space is ergodic (a function of energy alone).

Using this idea we are able to demonstrate that: (a) observational analyses that falsely assume spherical symmetry are made more accurate by imposing a strong prior preference for near-isotropic velocity dispersions in the centre of spheroids; (b) numerical simulations that use an idealised spherically-symmetric setup can yield misleading results and should be avoided where possible; and (c) triaxial dark matter halos (formed in collisionless cosmological simulations) nearly attain our maximally-stable limit, but their evolution freezes out before reaching it.

1 INTRODUCTION

Spherical symmetry is a foundational assumption of many dynamical analyses. The primary motivation is simplicity, since few astronomical objects are actually spherical. For example, observations and simulations both suggest that gravitational potential wells generated by dark matter halos are typically triaxial (e.g. Dubinski & Carlberg 1991; Cole & Lacey 1996; Jing & Suto 2002; Kasun & Evrard 2005; Hayashi et al. 2007; Schneider et al. 2012; Loebman et al. 2012). Characterising dark matter halos by spherically-averaged densities and velocities (e.g. Dubinski & Carlberg 1991; Navarro et al. 1996; Taylor & Navarro 2001; Stadel et al. 2009) at best tells only part of the story. At worst, it could be severely misleading.

The question of whether baryonic processes can convert dark matter cusps into cores (Pontzen & Governato 2014) provides one motivation for a detailed study of the relationship between spherical

and near-spherical dynamics. To explain why, we need to look forward to some of our results. Later in this paper, we cut a dark matter halo out of a cosmological simulation, then match it to an exactly spherical halo with an identical density and velocity anisotropy profile. This gives us two easy-to-compare equilibrium structures – the first triaxial, the second spherical – to perform a dynamical comparison. We expose each to the same time-varying gravitational potential, mimicking the effects of stellar feedback (there are no actual baryons in these runs). After 1 Gyr, the triaxial halo’s averaged density profile flattens into a convincing dark matter core, but the spherical halo maintains its cusp (see Figure 1).

This example, which is fully explored in Section 3.5, illustrates how it is dangerous to use spherically-symmetric simulations to infer anything about dynamics — even spherically-averaged dynamics — in the real universe. A spherical system does not evolve in the same way as the spherical averages of a triaxial system.

Ignoring asphericity can also lead to observational biases (e.g.

Hayashi & Navarro 2006; Corless & King 2007). From a dynamical standpoint, the nature of orbits in triaxial potentials is fundamentally different from those in spherical potentials: although the total angular momentum of any self-gravitating system must always be conserved, it is only in the spherical case that this conservation holds for individual particles. Conversely, a large fraction of dark matter particles near the centre of cosmological halos will be on box orbits which do not conserve their individual angular momenta (de Zeeuw & Merritt 1983; Merritt & Valluri 1996; Holley-Bockelmann et al. 2001; Adams et al. 2007). One practical consequence is that asphericity may be responsible for filling the loss cones of supermassive black holes at the centre of the corresponding galaxies (Merritt & Poon 2004).

Finally, it is known that asphericity plays a fundamental role in setting the equilibrium density profile during gravitational cold collapse (e.g. Huss et al. 1999). The underlying process is known as the radial orbit instability or ROI (Henon 1973; Saha 1991; Huss et al. 1999; MacMillan et al. 2006; Bellovary et al. 2008; Barnes et al. 2009; Maréchal & Perez 2009); a related effect was discussed by Adams et al. (2007). The name arises because particles on radial orbits are perturbed onto more circular trajectories. At the same time, the density distribution becomes triaxial. Even in the case of a uniform spherical collapse, this symmetry-breaking process is still triggered, presumably by numerical noise; the tangential component of forces must be unphysically suppressed for the system to remain spherical (Huss et al. 1999; MacMillan et al. 2006).

Despite all this, assuming spherical symmetry is very tempting because it makes life so much easier. Defining spherically-averaged quantities is a well-defined and sensible procedure even if we actually have the full distribution function in hand (as in simulations): departures from spherical symmetry are sufficiently small that different averaging procedures lead to consistent results (Saha & Read 2009). Additionally, when an aspherical halo is in equilibrium, we have shown numerically that a “sphericalised” version of it is also in equilibrium (see Appendix B of Pontzen & Governato 2013). This is helpful because it allows one to make a meaningful analysis in spherical coordinates, even when the system is aspherical. But it breaks down when out-of-equilibrium processes are included, as in the stellar-feedback-driven core-creation example above.

The present paper formalizes the idea of spherical analysis performed on aspherical systems and follows through the consequences. We will study equilibrium distribution functions in nearly, but not exactly, spherically-symmetric potentials, and focus on maximally-stable systems (which we define as being stable against all possible external linear perturbations). We will find that, in spherical coordinates, such systems appear to be “ergodic” (meaning that their distribution functions depend on energy alone) because the individual particles move randomly in angular momentum while maintaining a near-constant energy. It is important to emphasise that this describes the appearance of the system when analysed in a spherical coordinate system and the true system need not be chaotic for the result to hold, provided any isolating integrals are not closely related to angular momentum.

The formal statement of this idea is derived in Section 2. A brief overview of the required background is given in Appendix A, and a second-order derivation of the evolution is given in Appendix B. Section 3 outlines the practical consequences, starting by recasting and extending the phenomenology of the radial orbit instability. We describe an immediate implication for observational studies of aspherical systems, namely a new way to break the anisotropy degeneracy. Finally we return to the motivating problem above and explain why triaxial systems can undergo cusp-core tran-




Figure 1. One motivation for studying the relationship between spherical and aspherical dynamics is that the conversion of a dark matter cusp into a core by baryonic processes is qualitatively different in the two cases. Here we show the inner log density slope from a numerical experiment on two halos. One is spherical (dashed line) and the other triaxial (solid line) but their spherically-averaged properties are initially identical. An external potential has been added at the centre during the times indicated by the grey bands, with the fluctuations mimicking stellar feedback. The triaxial halo develops a clear core, whereas the spherical halo almost maintains its central density cusp. A complete description and analysis is given in Section 3.5.

sitions more easily than spherical systems. Section 4 concludes and points to open questions and future work.

2 ASPHERICAL DYNAMICS IN SPHERICAL COORDINATES

In this section we consider an aspherical system which is maximally stable against external linear perturbations. We assume that an observer of this system analyses it assuming spherical symmetry. We will show that this observer (falsely) concludes that the system is ergodic, *i.e.* that the density of particles in phase space is a function of energy alone. The derivation requires the use of action-angle coordinates; a crash course is provided in Appendix A.

2.1 Single particles

Given any near-spherical system, the Hamiltonian in the spherical action-angle variables is

$$H(\mathbf{J}, \boldsymbol{\Theta}) = H_0(\mathbf{J}) + \delta H(\mathbf{J}, \boldsymbol{\Theta}), \quad (1)$$

where H_0 is the sum of kinetic and potential energies in the spherical background, $\mathbf{J} = (J_r, j, j_z)$ is the vector of spherical actions (see Appendix A), $\boldsymbol{\Theta}$ is the vector of spherical angles and δH contains the perturbation (which includes the aspherical correction to the potential). The orbit of a particle in exact spherical symmetry, $\delta H = 0$, is described by Hamilton’s equations:

$$\dot{\mathbf{J}}_0 = -\frac{\partial H_0}{\partial \boldsymbol{\Theta}} = 0; \quad \dot{\boldsymbol{\Theta}}_0 = \frac{\partial H_0}{\partial \mathbf{J}} \Big|_{\mathbf{J}=\mathbf{J}_0} \equiv \boldsymbol{\Omega}_0(\mathbf{J}_0), \quad (2)$$

which defines the constant background orbital frequencies $\boldsymbol{\Omega}_0(\mathbf{J}_0)$. The expressions \mathbf{J}_0 , $\boldsymbol{\Theta}_0$ and $\boldsymbol{\Omega}_0(\mathbf{J}_0)$ will be used throughout to refer to the background ($\delta H = 0$) solution. This algebraically simple

form of the equations of motion is the reason for using action-angle variables, since it immediately integrates to

$$\mathbf{J}_0(t) = \mathbf{J}_0 = \text{constant}; \quad \Theta_0(t) = \Theta_0(0) + \Omega_0(\mathbf{J}_0)t, \quad (3)$$

where \mathbf{J}_0 and $\Theta_0(0)$ specify the initial action and angle coordinates of the orbit.

We now consider the effect of the aspherical correction to the potential encoded in δH , using standard Hamiltonian perturbation theory (e.g. Lichtenberg & Lieberman 1992; Binney & Tremaine 2008). First, taking advantage of the angle coordinates Θ being periodic in 2π , δH is expressed as

$$\delta H(\mathbf{J}, \Theta) = \sum_n \delta H_n(\mathbf{J}) e^{in \cdot \Theta}. \quad (4)$$

This equation states that, at any fixed \mathbf{J} , one can expand the periodic Θ dependence in a Fourier series without loss of generality.

We are interested in the evolution of \mathbf{J} at first order in the perturbation. Hamilton's relevant equation now reads:

$$\dot{\mathbf{J}} = -\frac{\partial H}{\partial \Theta} = -\sum_n in \delta H_n(\mathbf{J}) e^{in \cdot \Theta}. \quad (5)$$

Because δH is small, the result to first-order accuracy is given by substituting the zero-order solution (3) into equation (5) and integrating to give

$$\mathbf{J}(t) = \mathbf{J}_0 - \sum_n \frac{n}{\mathbf{n} \cdot \Omega_0} \delta H_n(\mathbf{J}_0) e^{in \cdot \Theta_0(t)} + \dots \quad (6)$$

Consequently as $\mathbf{n} \cdot \Omega_0 \rightarrow 0$, the linear-order correction to the orbit of a particle can become large even if the aspherical correction to the potential (δH) is small, an effect known as resonance (Binney & Tremaine 2008). Consider now the evolution of the background energy along the perturbed trajectory, given by

$$H_0(\mathbf{J}(t)) \simeq H_0(\mathbf{J}_0) + \frac{\partial H_0}{\partial \mathbf{J}} \cdot (\mathbf{J}(t) - \mathbf{J}_0) + \dots, \quad (7)$$

where we have Taylor-expanded to first order around \mathbf{J}_0 . Substituting equation 6 for $\mathbf{J}(t)$, there is a cancellation between numerator and denominator:

$$H_0(\mathbf{J}(t)) \simeq H_0(\mathbf{J}_0) - \sum_n \delta H_n(\mathbf{J}_0) e^{in \cdot \Theta_0(t)} + \dots \quad (8)$$

and so the fractional variation in H_0 remains small, even if \mathbf{J} changes significantly over time. In other words, according to linear perturbation theory, particles migrate large distances in \mathbf{J} along surfaces of approximately constant background energy H_0 . One can verify this constrained-migration prediction in numerical simulations of dark matter halos – an explicit demonstration is given in Appendix A3. This gives us some intuition for the result to come: a population of particles will seem to “randomise” their actions (in-closing angular momentum), but not their energy distribution.

The extent of the migration will depend on the nature of the potential in which a particle orbits. To quantify this requires going beyond linear perturbation theory and is the subject of “KAM theory” after Kolmogorov (1954), Arnold (1963) and Moser (1962); see e.g. Binney & Tremaine (2008); Goldstein et al. (2002); Lichtenberg & Lieberman (1992) for introductions. Colloquially the result is that for any given small perturbation the migration of typical orbits is also small. Arnold diffusion offers the most famous route to more significant diffusion through action space (see Lichtenberg & Lieberman 1992); but in our case, there is a more immediate reason why the KAM result does not in fact hold. Specifically, KAM theory relies on the frequencies $\Omega(\mathbf{J})$ being non-degenerate – i.e. that any change in the action leads also to a change in the frequencies, thus shutting off resonant migration. In smooth potentials, Ω

is almost a function of energy alone (see Appendix A3) and so the migration can be substantial.

Overall we informally expect particles to redistribute themselves randomly within the action shell of fixed background energy until they are evenly spread, implying a distribution function that appears ergodic in a spherical analysis. This does not imply the orbits are chaotic in the traditional sense; it is only because we are analysing an aspherical system in spherical coordinates that the phenomenon arises. With this in mind, we now turn to a more formal demonstration of the result.

2.2 The distribution function

So far we have discussed how a single particle orbiting in a mildly aspherical potential does not conserve its spherical actions (e.g. angular momentum). We informally suggested that a population will appear to ‘randomise’ the spherical actions at fixed energy. We now show more formally that a distribution function of particles subject to aspherical perturbations will be most stable when it is spread evenly on surfaces of constant H_0 .

We start by decomposing the true distribution function of particles in phase space, f , in terms of a spherical background f_0 and a perturbation δf . To make sure the split between spherical background and aspherical perturbation is uniquely defined, we take f_0 as the distribution function obtained when we perform a naive analysis averaging out the aspherical contribution:

$$f_0(\mathbf{J}) \equiv \frac{1}{(2\pi)^3} \int d^3\Theta f(\mathbf{J}, \Theta). \quad (9)$$

By Jean's Theorem, f_0 is an equilibrium distribution function in the spherical background because it is constructed from spherical invariants \mathbf{J} alone. Analogous to equations (1) and (4) one can write the full distribution function f as

$$\begin{aligned} f(\mathbf{J}, \Theta) &= f_0(\mathbf{J}) + \delta f(\mathbf{J}, \Theta) \\ &= f_0(\mathbf{J}) + \sum_n \delta f_n(\mathbf{J}) e^{in \cdot \Theta}. \end{aligned} \quad (10)$$

The whole f is to be in equilibrium in the true system, $\partial f / \partial t = 0$. We can turn this into an explicit condition on f using the collisionless Boltzmann equation,

$$0 = \frac{\partial f}{\partial t} = [H, f] \equiv \frac{\partial H}{\partial \Theta} \cdot \frac{\partial f}{\partial \mathbf{J}} - \frac{\partial H}{\partial \mathbf{J}} \cdot \frac{\partial f}{\partial \Theta}. \quad (11)$$

Expanding to linear order in H and f gives the condition

$$\sum_n \left(\Omega_0(\mathbf{J}) \cdot \mathbf{n} \delta f_n(\mathbf{J}) - \frac{\partial f_0}{\partial \mathbf{J}} \cdot \mathbf{n} \delta H_n(\mathbf{J}) \right) e^{in \cdot \Theta} = 0. \quad (12)$$

The different Θ dependence of each term in the sum means that the term in brackets must be zero for each \mathbf{n} . In particular, for the resonant terms \mathbf{n}_\perp where $\Omega_0 \cdot \mathbf{n}_\perp = 0$ one has the condition

$$\left(\frac{\partial f_0}{\partial \mathbf{J}} \cdot \mathbf{n}_\perp \right) \delta H_{\mathbf{n}_\perp}(\mathbf{J}) = 0. \quad (13)$$

A sufficient condition for stability of δf is therefore that f_0 is a function only of H_0 , since then $\partial f_0 / \partial \mathbf{J} = \Omega_0 d f_0 / d H_0$ and consequently the dot product in equation (13) vanishes.

This is the core result claimed at the start of the section: $f_0 = f_0(H_0)$, i.e. the distribution function implied by a spherical analysis appears to be ergodic. It is not a necessary condition for achieving equilibrium, since for any given aspherical system certain $\delta H_{\mathbf{n}_\perp}$ will be zero. Rather, the result should be read as applying to the maximally stable distribution function – a distribution function that does not evolve under any linear perturbation to its potential.

We again emphasise that the distribution function f_0 is fictional. There is no sense in which the true distribution function, f , is actually ergodic. The statement is about how the system appears to be when it is analysed using spherically averaged quantities, equation (9). Yet, it establishes a way in which we can understand these spherically-averaged quantities in a systematic way – the system is most stable if it *appears* ergodic, regardless of what the underlying dynamics is really up to. In the remainder of this paper we will refer to such systems as ‘spherically ergodic’.

2.3 Testable predictions

We have established that systems which appear to be ergodic in a spherical analysis are maximally stable. Now we need to devise a connection to observable or numerically-measurable quantities.

A distribution function $f(H_0)$ that is truly a function of energy alone has an isotropic velocity distribution (Binney & Tremaine 2008). To test for isotropy, one calculates $\beta(r)$ according to the usual spherically-averaged definition

$$\beta(r) = 1 - \frac{\langle v_r^2 \rangle(r)}{2\langle v_t^2 \rangle(r)} \quad (14)$$

where v_r is a particle’s radial velocity, v_t its tangential velocity and the angle-bracket averages are taken in radial shells. For a population on radial orbits, $\beta(r) = 1$; conversely for purely circular motion, $\beta(r) = -\infty$. Between these two extremes, an ergodic population has $\beta(r) = 0$ (Binney & Tremaine 2008).

Intuitively, a spherically ergodic system (in the sense defined in the previous section) should therefore be approximately isotropic. However one has to handle that expectation with a little care because the true population f is *not* ergodic and the measured velocity dispersions, even in spherical polar coordinates, may inherit information from f that is not present in f_0 .

Instead we will now construct a more rigorously justifiable, slightly different statement that still connects spherically ergodic populations to velocity isotropy. Measuring the mean of any function of the spherical actions $q(\mathbf{J})$, we obtain

$$\int d^3J d^3\Theta f(\mathbf{J}, \mathbf{\Theta}) q(\mathbf{J}) = \int d^3J f_0(H_0(\mathbf{J})) q(\mathbf{J}), \quad (15)$$

an exact result. Therefore any statement about averages over spherical actions automatically knows only about f_0 – the spherical part of the distribution function. This allows us to derive unambiguous implications of a spherically ergodic population.

The most familiar action is the specific scalar angular momentum j . Because it is a scalar for each particle, averages over this quantity do not express anything about a net spin of the halo but rather about the mix of circular and radial orbits, just like the traditional velocity anisotropy. Radially-biased populations have $\langle j \rangle \simeq 0$ whereas populations on circular orbits have $\langle j \rangle = j_c$, where j_c is the maximum angular momentum available at a given energy. So, velocity anisotropy can be conveniently represented in terms of the mean scalar angular momentum.

We can go further and calculate a function, $\langle j \rangle(E)$, where the average is taken only over particles at a particular specific energy. This quantity can be represented in terms of the ratio of two integrals of the form (15):

$$\langle j \rangle(E) = \frac{\iiint dJ_r dJ_z f_0(H_0) j \delta(H_0 - E)}{\iiint dJ_r dJ_z f_0(H_0) \delta(H_0 - E)}. \quad (16)$$

The triple integral ranges over the physical phase space coordinates: $0 \leq J_r < \infty$, $0 \leq j < \infty$, $-j \leq J_z \leq j$. One can immediately

perform the J_z integrals; then the J_r integral can be completed by changing variables to H_0 (recalling $\partial H_0 / \partial J_r \equiv \Omega_r$) and consuming the δ function. After this manipulation j can only range between 0 and $j_c(E)$ where $j_c(E)$ is the specific angular momentum corresponding to a circular orbit with specific energy E ; there are no physical orbits with more angular momentum at the specified E . The final, exact result is:

$$\langle j \rangle(E) = \int_0^{j_c(E)} dj \Omega_r(E, j)^{-1} j^2 \Big/ \int_0^{j_c(E)} dj \Omega_r(E, j)^{-1} j. \quad (17)$$

We now have a firm prediction for spherically ergodic populations. Namely, if we bin particles in E and measure $\langle j \rangle(E)$ in each bin, the results should be predicted by equation 17, which is a function only of the potential (through Ω_r). Equation (17) does not exhaust the possible tests for spherical ergodicity, but it is sufficient for our present exploratory purposes.

For smooth potentials, $\Omega_r(E, j)$ varies very little between $j = 0$ and $j = j_c$, and one can approximate it very well as a function of E alone (Appendix A3). In this case, the integrals follow analytically and one has the result

$$\langle j \rangle(E) \simeq \frac{2}{3} j_c(E). \quad (18)$$

This is a helpful simplification to set expectations, but throughout this paper when showing the spherically ergodic limit, we will use the exact expression given by equation (17).

3 EXAMPLE APPLICATIONS

So far we have motivated and derived a formal result that aspherical systems are most stable when they appear ergodic in spherical coordinates. We derived one practical consequence for the angular momentum distribution, equation (17), which in an approximate sense states that the velocity distribution will appear isotropic. We are now in a position to test whether numerical simulations actually tend towards this maximally-stable limit in a variety of situations, beginning with cosmological collisionless dark matter halos.

3.1 Cosmological dark matter halos

Let us re-examine the three high-resolution, dark-matter-only zoom cosmological simulations used in the analysis of Pontzen & Governato (2013). The three each have several million particles in their $z = 0$ halos which correspond in turn to a dwarf irregular, L_* galaxy and cluster. The force softening ϵ , virial radius r_{200} (at which the mean density enclosed is 200 times the critical density) and virial masses M_{200} are 65, 170, 690 pc; 98, 301, 1430 kpc and 2.8×10^{10} , 8.0×10^{11} , $8.7 \times 10^{13} M_\odot$ respectively. For further details see Pontzen & Governato (2013).

The upper panel of Figure 2 shows the anisotropy for our cosmological halos. To compare the three directly, we scale the radius by r_{\max} (respectively 27, 57 and 340 kpc), the radius at which the circular velocity $(GM(< r)/r)^{1/2}$ reaches its maximum, v_{\max} ($= 56, 150$ and 610 km s^{-1}). We restrict attention to the region well within the virial radius; here, the anisotropy $\beta(r)$ typically lies between the purely radial and isotropic cases (e.g. Bellovary et al. 2008; Navarro et al. 2010).

We now want to link this relatively familiar velocity anisotropy to the alternative $\langle j \rangle(E)$ statistic that was directly predicted by the spherically ergodic property in Section 2. For each particle we calculate the specific energy $E = \dot{\mathbf{x}}^2/2 + \Phi(|\mathbf{x}|)$, where




Figure 2. The velocity anisotropy of the inner parts of three sample high-resolution cosmological dark matter halos (simulated without baryons), plotted as a function of radius (upper panel) and in energy shells (lower panel). The upper panel shows the classic velocity anisotropy $\beta(r)$ defined in the text, for which a purely radial population has $\beta = 1$ and a population on circular orbits $\beta = -\infty$. The lower panel shows our alternative in energy space which can be more precisely related to the theoretical arguments presented in Section 2; here $\langle j \rangle(E)/j_{\text{circ}} = 0$ for radial orbits and 1 for circular orbits, and approximately $2/3$ for an isotropic population. Both panels show that the halos have near-isotropic orbits with a slight radial bias. The range of the two plots is roughly comparable, but we caution that the mapping from r to E is not unique (see Figure 3).

\mathbf{x} is the vector displacement from the halo centre. To make this quantity agree exactly with H_0 in the terminology of Section 2, we ignore asphericity when calculating the potential energy, defining it as

$$\Phi(r) \equiv \int_0^r \frac{GM(<r')}{r'^2} dr' \quad (19)$$

where $M(<r')$ is the mass enclosed inside a sphere of radius r' . (The numerical integration is performed by binning particles in shells of fixed width ϵ , chosen to coincide with the force softening in the simulation; within these bins the density is taken to be constant.) The physics is invariant if a constant is added to the potential; we have chosen to fix its scale by setting $\Phi(0) = 0$.

For each particle we also calculate the specific angular momentum $j = |\mathbf{x} \times \dot{\mathbf{x}}|$, and the specific angular momentum of a circular orbit at the same energy, $j_{\text{circ}}(E)$ which is given by simultaneously solving

$$E = \Phi(r) + \frac{j_{\text{circ}}^2}{2r^2} \quad \text{and} \quad \Phi'(r) = \frac{j_{\text{circ}}^2}{r^3}, \quad (20)$$

to eliminate r in favour of j_{circ} .

We plot $\langle j \rangle(E)/j_{\text{circ}}(E)$ in bins containing 1000 particles each in the lower panel of Figure 2. To facilitate comparison with the top panel, a population on purely radial orbits would have $\beta = 1$ and $\langle j \rangle/j_{\text{circ}} = 0$, whereas a purely circular distribution function corresponds to $\beta = -\infty$ or $\langle j \rangle/j_{\text{circ}} = 1$. Isotropic, purely spherical populations have $\beta(r) = 0$ and $\langle j \rangle/j_{\text{circ}} \simeq 2/3$ as discussed at the end of Section 2.3. When compared against each other in this way, the two panels agree well: the populations are on near-isotropic orbits with a slight radial bias.

Quantitatively, how well do these results agree? As a guideline, we can compare results for models with constant anisotropy.




Figure 3. The relationship between r/r_{max} and E/v_{max}^2 for circular orbits (dashed line), radial orbits at apocentre (solid line) and particles in the ‘MW’ run (density shows the number of particles with energy E at each radius r). The mapping between energy and radius is fuzzy, so that anisotropy at high E can easily contaminate $\beta(r)$ at small r .

From Binney & Tremaine (2008), a distribution function $f(j, E) = j^{-2\beta} f_1(E)$ generates a constant anisotropy $\beta(r) = \beta$. We can calculate the connection to the new statistic by generalising the reasoning of Section 2.3, with the result that

$$\langle j \rangle(E) = \frac{\int_0^{j_{\text{circ}}} dj \Omega_r^{-1} j^{2-2\beta}}{\int_0^{j_{\text{circ}}} dj \Omega_r^{-1} j^{1-2\beta}} \simeq \frac{2-2\beta}{3-2\beta} j_{\text{circ}}, \quad (21)$$

where the first result is exact and the second follows from assuming Ω_r is independent of j (which is an excellent approximation; see Appendix A3). Consistent with equation (18), $\beta = 0$ gives $\langle j \rangle(E)/j_{\text{circ}} \simeq 2/3$. But as the system becomes more radially biased we can now calculate that, for example, $\beta = 0.2$ corresponds to $\langle j \rangle(E)/j_{\text{circ}} \simeq 0.62$. Despite the various approximations involved, these values therefore correctly relate the values of β in the top panel of Figure 2 with the $\langle j \rangle$ results in the lower panel.

That said, a detailed comparison as a function of radius is hard because particles at a given radius r have a wide spread of energies E . Figure 3 illustrates this relationship for the ‘MW’ halo. The density shows the probability of a particle at radius r also having specific energy E , $p(E|r)$. The minimum E at each r is set by $\Phi(r)$, which gives the energy of a particle at apocentre (solid line). A more typical E is given by the energy of a circular orbit at r (dashed line), and this gives some intuition for mapping results from the top panel of Figure 2 onto the bottom panel. However, any E exceeding $\Phi(r)$ is theoretically possible. So β at any given radius actually represents an average over particles of many different energies.

3.2 The classical radial orbit instability

Having established that, loosely speaking, $\langle j \rangle(E)$ represents the anisotropy in energy shells in the same way that $\beta(r)$ does in radial shells, we can return to our prediction (Section 2.3) for the former quantity, which is shown by the dashed line in the lower panel of Figure 2. The prediction is almost, but not quite, satisfied in cos-




Figure 4. The $\beta(r)$ velocity anisotropy (upper panel) and its equivalent in energy space $\langle j \rangle(E)$ (lower panel, as Figure 2) for the evolution over time of a halo that is initially in equilibrium, but unstable against the radial orbit instability. The initial conditions are represented by the dark blue line, $t = 0$; there is a delay of $12t_{\text{dyn}}$ before any significant evolution can be seen (light blue line). Subsequent lines are shown every $4t_{\text{dyn}}$. Once the instability kicks in and generates aspherical perturbations, there is a rapid evolution towards the predicted spherically ergodic limit (lower panel).

mological dark matter halos. Since the condition is only reached in a maximally-stable object, approximate agreement is an acceptable situation.

Because cosmological halos initially form from near-cold collapse, the radial orbit instability (van Albada 1982; Barnes et al. 1986; Saha 1991; Weinberg 1991) is invoked to explain how the radially infalling material gets scattered onto a wider variety of orbits MacMillan et al. (2006); Bellovary et al. (2008); Barnes et al. (2009), isotropising the velocity dispersion. Most tellingly, numerical experiments by Huss et al. (1999); MacMillan et al. (2006) show that the suppressing the instability (by switching off non-radial forces) results in a qualitatively different density profile as the end-point of collapse. We will now show that the isotropisation of velocity dispersion during the radial orbit instability can be interpreted in terms of an evolution towards stability in the terms of Section 2.

Consider what happens to a halo that is intentionally designed to be unstable. First, we will show the classic radial orbit instability (ROI) at work by constructing a spherical halo with particles that are on radially-biased orbits. We initialise our particles such that they solve the Boltzmann equation and so are stable in exact spherical symmetry. In practice, however, the strong radial bias means that any slight numerical noise will trigger the ROI. By initialising an unstable equilibrium in this way, we avoid confusion from violent relaxation processes associated with out-of-equilibrium collapse (Lynden-Bell 1967).

Specifically, the initial conditions are set up in a similar fashion to Read et al. (2006), with particle positions drawn from a generalised Hernquist density profile (Hernquist 1990; Dehnen 1993):

$$\rho(r) = \frac{M(3-\gamma)}{4\pi a^3} \left(\frac{r}{a}\right)^{-\gamma} \left(1 + \frac{r}{a}\right)^{\gamma-4}, \quad (22)$$

which has a circular velocity reaching a peak at $r_{\text{max}} = (2-\gamma)a$ and implies the gravitational potential

$$\Phi(r) = \frac{GM}{2-\gamma} \left[\left(1 + \frac{r}{a}\right)^{\gamma-2} - 1 \right]. \quad (23)$$

We choose $\gamma = 1$ to roughly mimic an NFW halo (Navarro et al. 1997) in the innermost parts of interest. The velocities are sampled (using an accept-reject algorithm) from a numerically-calculated Osipkov-Merritt distribution function (Binney & Tremaine 2008):

$$f(Q) = \frac{\sqrt{2}}{4\pi^2} \frac{d}{dQ} \int_Q^0 \frac{d\Phi}{\sqrt{\Phi-Q}} \frac{d}{d\Phi} \left[\left(1 + r(\Phi)^2/r_a^2\right) \rho(r(\Phi)) \right] \quad (24)$$

where the parameter Q is defined by

$$Q \equiv E + \frac{j^2}{2r_a^2}. \quad (25)$$

The value of r_a is known as the anisotropy radius because the velocity anisotropy is given by

$$\beta(r) = \frac{1}{1 + r_a^2/r^2}, \quad (26)$$

showing that $\beta(r) \simeq 0$ (isotropic) for $r \ll r_a$ and $\beta(r) \simeq 1$ (radial) for $r \gg r_a$. We used the minimum value of r_a for which the distribution function is everywhere positive, making the orbits as radially biased as possible; for $\gamma = 1$, this is $r_a \simeq 0.21$ (Meza & Zamorano 1997). We draw 10^6 particles and evolve the system using RAMSES (Teyssier 2002), with mesh refinement based on the number of particles per cell, resulting in a naturally adaptive force softening reaching a minimum of $\varepsilon = 90\text{pc}$.

Our expectation that numerical noise triggers the ROI is borne out by the numerical experiment. The upper panel of Figure 4 shows the radial anisotropy $\beta(r)$ over time. We have defined a single dynamical time, t_{dyn} , at the peak of the velocity curve so that $t_{\text{dyn}} \equiv r_{\text{max}}/v_{\text{max}}$. The six solid lines show the population at $t = 0, 12, 16, 20, 24$ and 28 dynamical times. At first, $\beta(r)$ appears stable, but suddenly after $16t_{\text{dyn}}$ it becomes considerably more isotropic. Over the same timescale, asphericity in the potential develops. To demonstrate this, we determine the inertia tensor of the entire density distribution and calculate the ratio of the principal axes; at $t = 0$, the ratio is 1.0 by construction. By $16t_{\text{dyn}}$ the ratio is 1.8. It stabilises at around $26t_{\text{dyn}}$ with a ratio of 4.3.

This is symptomatic of the classic radial orbit instability in action. We can follow the same process from our energy/angular-momentum standpoint in the bottom panel of Figure 4. The initial conditions ($\Delta t = 0$) have a kink in them, with most regions of energy space appearing radially biased but some showing a slight circular preference (around $E \simeq 1.5v_{\text{max}}^2$). We verified that this is an artefact of the Osipkov-Merritt construction and is consistent with the radial $\beta(r)$ being radially biased everywhere; recall Figure 3 shows that the relationship between r and E is non-trivial.

As time progresses, the $\langle j \rangle(E)$ distribution correctly tends towards the spherically ergodic (SE) limit, as predicted. The SE limit is attained to good accuracy for energies $E < v_{\text{max}}^2$ by $\Delta t = 20t_{\text{dyn}}$. At larger energies, it is likely difficult to achieve SE because of the long time-scales and weak gravitational fields involved. Like $\langle j \rangle$, $\beta(r)$ at $28t_{\text{dyn}}$ shows an isotropic distribution at the centre ($r \rightarrow 0$); but $\beta(r) > 0$ everywhere for $r > 0.1 r_{\text{max}}$. We verified that this continuing radial bias throughout the halo is produced by a number of high-energy, loosely-bound particles plunging through, i.e. the high- E particles from the lower panel in Figure 4 pervade all radii in the upper panel.

Let us briefly recap: what we have so far is a new view of an existing phenomenon. We have recast the ROI and its impact on spherically-averaged quantities as an evolution towards an analytically-derived maximally stable class of distribution functions. Viewed in energy shells instead of radial shells, the distinction between the regions that reach the SE limit and the slowly-




Figure 5. The mean angular momentum of particles as a function of energy, like the lower panel of Figure 4 – but now for a subpopulation in a dark matter halo that is globally in equilibrium. At $\Delta t = 0$, we select particles on predominantly radial ($j/j_{\text{circ}} < 0.2$) orbits; at later times, the subpopulation mean evolves back towards the population mean. The lines from bottom to top show the state at selection and after 0.2, 0.4, 0.6, 1.0 and $4.0t_{\text{dyn}}$ respectively.

evolving, loosely-bound regions that remain radially biased is considerably clearer. Now, because our underlying analytic result is derived without requiring the population to be self-gravitating, we can now go further and consider a different version of the radial orbit instability – and show that it continues to operate even when a system is in a completely stable equilibrium.

3.3 The continuous radial orbit instability

Our next target for investigation is to broaden the conditions: according to the results of Section 2, a subpopulation of particles should undergo something much like the radial orbit instability even when the *global* potential is completely stable. A suitable name for this phenomenon would seem to be the “continuous radial orbit instability”, since it continues indefinitely after the global potential has stabilised.

We perform another numerical experiment to demonstrate the effect. First, to avoid confusion from cosmological infall and tidal fields, we create a stable, isolated, triaxial cosmological halo by extracting from our cosmological run a region of $3r_{200}$ around our ‘Dwarf’. We then evolve this isolated region for 2 Gyr to allow any edge effects to die away, and verify that the density profile out to r_{200} is completely stable. As before we define a dynamical time for the system of $t_{\text{dyn}} \equiv r_{\text{max}}/v_{\text{max}} = 470 \text{ Myr}$.

After the 2 Gyr $\simeq 4t_{\text{dyn}}$ has elapsed, we select all particles with $j < 0.2j_{\text{circ}}$. These particles are, at the moment of selection, on preferentially radial orbits. We trace our particles forward through time, measuring $\langle j \rangle(E)$ in each snapshot. The results are shown in Figure 5 for various times between $\Delta t = 0.0$ and $4.0t_{\text{dyn}}$ after selection. Over this period, the mean angular momentum significantly increases towards the spherically ergodic limit at every energy. The changes are much faster at low energies where the particles are tightly bound and the local dynamical time is short compared to the globally-defined t_{dyn} .

The evolution is rapid until $4t_{\text{dyn}}$ after which the $\langle j \rangle(E)$ remains near-static (except at $E > 2.5v_{\text{max}}^2$ where it continues to




Figure 6. The $\beta(r)$ (upper panel) and $\langle j \rangle(E)$ (lower panel) relations for the stellar populations of dwarf satellites around a Milky-Way-like central. In the upper panel arrows indicate the location of $r_{1/2}$, the radius enclosing half the stellar mass. The profile is plotted exterior to 3ϵ where ϵ is the softening length. Because of the tidal interaction with the parent galaxy, the outer parts of each object are out of equilibrium and display biases ranging from strongly radial to circular. The lower panel shows the same populations in energy space. At low energies, tightly bound particles are now seen to be close to the spherically ergodic limit. At higher energies, a spread is seen but not as large as would naively be expected from the $\beta(r)$ relation.

slowly rise). Eventually the subpopulation has $\langle j \rangle(E) \simeq 0.5j_{\text{circ}}(E)$ independent of E . This establishes that particles at low angular momentum within a completely stable, unevolving halo, automatically evolve towards higher angular momentum. After a few dynamical times, their angular momentum becomes comparable to that of a randomly-selected particle from the full population, although with a continuing slight radial bias.

One can understand this incompleteness in a number of equivalent ways. Within our formal picture, it arises from the fact that only certain δH_n are non-zero in equation (12). More intuitively, the continuing process of subpopulation evolution is being driven by particles on box orbits that change their angular momentum at near-fixed energy – but not all particles are on such orbits, and so some memory of the initial selection persists.

3.4 Observations of dwarf spheroidal galaxies

In the previous section, we established that subpopulations on initially radially-biased orbits evolve towards velocity isotropy even if the global potential is stable. We explained this in terms of our earlier calculations, and now turn to how the results might impact on observations.

An important current astrophysical question is how dark matter is distributed within low-mass galaxies: this can discriminate between different particle physics scenarios (Pontzen & Governato 2014). The smallest known galaxies, dwarf spheroidals surrounding our Milky Way, in principle provide a unique laboratory from this perspective. But determining the distribution of dark matter is a degenerate problem because of the unknown transverse velocities of the stellar component (e.g. Charbonnier et al. 2011). Furthermore the use of spherical analyses seems inappropriate since the underlying systems are known to be triaxial (e.g. Bonnavard et al. 2015).

This is exactly the kind of situation we set up in our initial calculations (Section 2): an aspherical system being analysed in spherical symmetry. We have shown that the results apply both to tracer and self-gravitating populations. So, we can go ahead and apply it to the stars in observed systems: the stellar population of a dwarf spheroidal system will be maximally stable if it appears ergodic in a spherical analysis. This implies a strong prior on what spherical distribution functions are actually acceptable and therefore, in principle, lessens the anisotropy degeneracy.

This idea warrants exploration in a separate paper; here we will briefly test whether the idea is feasible by inspecting some simulated dwarf spheroidal satellite systems. In particular we use a gas-dynamical simulation of “MW” region (see above) using exactly the same resolution and physics as Zolotov et al. (2012); see that work for technical details (although note the actual simulation box is a different realisation).

We analyse all satellites with more than 10^4 star particles, which gives us objects lying in the ranges $5.4 \times 10^7 < M_\star < 4.9 \times 10^8 M_\odot$, $28 < v_{\max} < 60 \text{ km s}^{-1}$ and $4.6 < r_{\max} < 9.6 \text{ kpc}$. We first verified that these do not host rotationally supported disks; we then calculated $\beta(r)$ profiles between $3\varepsilon \simeq 0.5 \text{ kpc}$ and $0.6 r_{\max}$. The stellar half-light radius is $0.15 < r_{1/2}/r_{\max} < 0.3$, so our calculations extend into the outer edges of each visible system.

The results are shown in the upper panel of Figure 6, with arrows indicating $r_{1/2}$. The outer regions display a range of different behaviours from strongly radial to circular. However, in each case the $\beta(r)$ is nearly isotropic as $r \rightarrow 0$. This is consistent with a view in which the centres have achieved stability while the outer parts are being harassed by the parent tides and stripping.

The picture is reinforced by considering the $\langle j \rangle(E)$ statistic (lower panel of Figure 6). Here, stars on tightly bound orbits (small E) are very close to the spherically ergodic limit. In the less-bound regions, the agreement worsens. However a quantitative analysis using the approximate relation (21) shows that the $\langle j \rangle(E)$ relation stays much closer to the spherically ergodic limit than the $\beta(r)$ relation naively implies. This is probably because the shape of $\beta(r)$ is in part determined by out-of-equilibrium processes, coupled to the multivalued relationship between r and E (Figure 3).

In any case these results suggest that we should adopt priors on spherical Jeans or Schwarzschild analyses that strongly favour near-isotropy in the centre of spheroidal systems – or, more accurately, strongly favour near-ergodicity for tightly bound stars. In future work we will expand on these ideas and apply them to observational data, since they could lead to a substantial weakening of the problematic density-anisotropy degeneracies.

3.5 Dark matter cusp – core transitions

In Pontzen & Governato (2012, henceforth PG12) we established that dark matter can be redistributed when intense, short, repeated bursts of star formation repeatedly clear the central regions of a forming dwarf galaxy of dense gas (see also Read & Gilmore 2005; Mashchenko et al. 2006). The resulting time-changing potential imparts net energy to the dark matter in accordance with an impulsive analytic approximation. This type of activity has now been seen or mimicked in a large number of simulations, allowing glimpses of the dependency of the process on galactic mass, feedback type and efficiency (Governato et al. 2012; Peñarrubia et al. 2012; Martizzi et al. 2012; Teyssier et al. 2013; Zolotov et al. 2012; Garrison-Kimmel et al. 2013; Di Cintio et al. 2014; Oñorbe et al. 2015).

We based our PG12 analysis on 3D zoom cosmological simulations, but our analytic model assumed exact spherical symmetry.




Figure 7. The velocity anisotropy ($\beta(r)$, upper panel) and angular momentum ($\langle j \rangle(E)$, lower panel) for cusped and cored halos from PG12; note the much-expanded y-axis scales relative to previous figures, which are required to highlight the differences. The cored cases (solid lines) are almost perfectly spherically ergodic (lower panel) and hence isotropic (upper panel). This contrasts with the cusped case (dotted lines) which has a slight but significant radial bias (seen as high β in the upper panel and low j in the lower panel).

By definition, in the analytic model all particles conserve their angular momentum at all times. Taken at face value, energy gains coupled to constant angular momentum would leave a radially biased population in the centre of the halo.

The top panel of Figure 7 shows the measured velocity anisotropy in the inner 5 kpc for the zoom simulations in PG12 at $z = 0$; the solid line shows the feedback simulation which has developed a core, whereas the dotted line shows the dark-matter-only simulation which maintains its cusp. The difference between the two cases is the opposite to that naively expected: the centre of the cored halo has a more isotropic velocity dispersion than the centre of the cusped halo. The lower panel shows the equivalent picture in energy space (noting that $\Phi(0) = 0$ and $\Phi(5 \text{ kpc}) = 3790 \text{ km}^2 \text{ s}^{-2}$). For clarity only a small fraction of the $0 < j/j_{\text{circ}} < 1$ interval is now plotted; the differences between the cusped (dotted) and cored (solid) lines are relatively small, but significant. At all energies the cored profile has more angular momentum than the cusped profile; in fact, it lies right on the spherically-ergodic limit (dashed line) whereas the cusped profile is biased by around 0.06 to radial orbits¹.

While the PG12 model correctly describes the energy shifts, it misses the angular momentum aspect of core creation, which has been emphasised elsewhere (e.g. Tonini et al. 2006). We will now show that a complete description of cusp–core transitions involves two components: energy gains consistent with PG12, and re-stabilisation of the distribution function consistent with the description of Sections 2.2 and 3.3. We again use the RAMSES code to follow isolated halos with particle mass $2 \times 10^4 M_\odot$; however we adapted the code to add an external, time-varying potential to the self-gravity.

¹ The dashed line for the isotropic j/j_{circ} in the lower panel is calculated using the cored potential. However the dependence on potential is very weak, as discussed in Section 3; calculating it instead with the cusped potential leads to differences of less than one percent at every energy.




Figure 8. The PG12 mechanism is reproduced by adding an external potential (to mimic ‘gas’) to a DM-only simulation for the time intervals indicated by the grey bands. We measure the response of a completely spherical halo (dotted lines) and a realistically triaxial halo (solid lines) as described in the text. In both cases we track the mean energy (top panel) and angular momentum (middle panel) of the 0.1% of particles that start out most bound. We also measure the log slope of the density profile at each time output (bottom panel). The energy shift (top) predicted by PG12 is found to hold in both cases. In the spherical case, the angular momentum remains constant up to discreteness effects; it therefore drops relative to the circular angular momentum (middle panel). Conversely, in the triaxial case, the continuous radial orbit instability causes the angular momentum to rise in proportion to j_{circ} . Only when the angular momentum is allowed to rise does a significant core form, meaning that the spherical case unphysically suppresses core formation (lower panel).

We took our stabilised, isolated, triaxial dark matter halo extracted as described in Section 3.3 and created a sphericalised version of it as follows. For each particle we generate a random rotation matrix following an algorithm given by Kirk (1992), then multiply the velocity and position vector by this matrix. Finally, we verified that the final particles are distributed evenly in solid angle, and that the spherically-averaged density profile and velocity anisotropy is unchanged. The triaxial and spherical halos are both NFW-like and stable over more than 2 Gyr when no external potential is applied.

For our science runs, we impose an external potential corresponding to $10^8 M_\odot$ gas in a spherical ball 1 kpc in radius, distributed following $\rho \propto r^{-2}$; this implies a potential perturbation of $\Delta\Phi = -700 \text{ km}^2 \text{s}^{-2}$ at 500 pc, for instance. The potential instantaneously switches off at 100 Myr, back on at 200 Myr, off at 300 Myr and so forth until it has accomplished four “bursts”. The period and the mass in gas is motivated by Figures 1 and 2 of PG12 and Figure 7 of T+13. We also tried imposing potentials with different regular periods and with random fluctuations, none of which altered the behaviour described below.

The top two panels of Figure 8 show the time-dependent behaviour of the central, most-bound 0.1% of particles. The triaxial and spherical halo results are illustrated by solid and dashed lines respectively. The top panel shows how the coupling of the external potential is similar; in particular, it results in a mean increase in specific energy of $\Delta E \simeq 200 \text{ km}^2 \text{s}^{-2}$ for particles in both cases. The final shift in the spherical case is very slightly larger than that in the triaxial case, a difference which is unimportant for what follows (it would tend to create a larger core if anything).

The middle panel displays the mean specific angular momentum j of the same particles as a fraction of the circular angular momentum $j_{\text{circ}}(E)$. In the spherical case (dashed line), this quantity drops significantly because j is fixed but j_{circ} is rising over time (since it increases with E). By contrast in the triaxial case j/j_{circ} returns to its original value, meaning that j has risen by the same fraction as $j_{\text{circ}}(E)$. This is dynamic confirmation of the discussion earlier in this Section: the spherically symmetric approach predicts an increasing bias to radial orbits – whereas in the realistic aspherical case, the stability requirements quickly erase this bias.

The bottom panel of Figure 8 shows the measured density slope $\alpha = d \ln \rho / d \ln r$ at 500 pc for the two experiments. Both start at $\alpha \simeq -1.5$; the triaxial case correctly develops a core (with $\alpha \simeq -0.1$) whereas the completely spherical case maintains a cusp (with $\alpha \simeq -0.9$). This discrepancy is causally connected to the angular momentum behaviour: the mean radius of a particle increases with increasing angular momentum j , even at fixed energy E , so a typical particle migrates further outwards in the triaxial case compared to the spherical.

We can therefore conclude that asphericity is a pre-requisite for efficient cusp–core transitions. However there is a subsidiary issue worth mentioning. After about 3 Gyr we find that the spherical halo autonomously does start increasing $\langle j \rangle / j_{\text{circ}}$ and the dark matter density slope becomes shallower. This is because the potential fluctuations have generated a radially biased population which is unstable, and a global radial orbit instability is triggered by numerical noise over a sufficiently long period (as in Section 3.2).

In fact, with ‘live’ baryons rather than imposed fluctuations, there are aspherical perturbations which accelerate the re-equilibration process further and renew a global radial orbit instability (Section 3.2), encouraging the entire population towards spherical ergodicity. As we saw in Figure 7, the cored halo from PG12 has an almost perfectly spherically-ergodic population for $E < 3000 \text{ km}^2 \text{s}^{-2}$ (unlike the cusped case). We verified that this is also true of T+13. Generating a core through potential fluctuations seems to complete a relaxation process that otherwise freezes out at an incomplete stage during collisionless collapse.

4 CONCLUSIONS

Let us return to the original question: how much do inaccuracies inherent in the spherical approximation really matter in practical situations? The answer is, unfortunately, that “it depends”; but we can now distinguish two cases as a rule of thumb:

- 1) For dynamical calculations or simulations, the inaccuracy matters a great deal. Neglecting the aspherical part of the potential unphysically freezes out the radial orbit instability and related effects, so can lead to qualitatively incorrect behaviour.
- 2) For the *analysis* of observations or simulations in equilibrium, the assumption is far more benign – it is, in fact, extremely powerful when handled with care. The underlying aspherical system and the fictional spherical system both appear to be in equilibrium; the mapping between the two views yields striking insight into (for example) spheroidal stellar distribution functions and dark matter halo equilibria.

These conclusions are based on the fact that, when aspherical systems are analysed in spherical coordinates, there is an attractor solution for the spherically-averaged distribution function f_0 – namely, it tends towards being ergodic (i.e. f_0 is well-approximated as a function of energy alone). We demonstrated this using Hamil-

tonian perturbation theory (see Section 2 and Appendix B), and subsequently used the term “spherically ergodic” (SE) to describe a distribution function f with the property that its spherical average f_0 is ergodic in this way.

The result follows because the orbits do not respect spherical integrals-of-motion such as angular momentum. Note, however, that the physical orbits may still possess invariants in a more appropriate set of coordinates. The apparent chaotic behaviour and tendency of particles to spread evenly at each energy is a helpful illusion caused by adopting coordinates that are only partially appropriate.

In Section 3 we showed that the idea has significant explanatory power. First, we inspected the selection of equilibrium in triaxial dark matter halos (Section 3.1) which led us to consider the classical radial orbit instability (Section 3.2). We demonstrated that the instability naturally terminates very near our SE limit (lower panel, Figure 4). Particles at high energies have long dynamical times which causes them to freeze out: they evolve towards, but do not reach, the SE limit. Because these high-energy particles often stray into the innermost regions (Figure 3), the velocity anisotropy $\beta(r)$ continues to display a significant radial bias at all r after the instability has frozen out. Moreover in the case of self-consistently formed cosmological halos (Figure 2), even at low energies there is a slight radial bias. The bias is only erased when a suitable external potential is applied (e.g. during baryonic cusp-core transformations), forcing the system to stabilise itself against a wider class of perturbations than it can self-consistently generate (Figure 7); we will return to this issue momentarily.

One novel aspect of our analysis compared to previous treatments of the radial orbit instability is that it applies as much to tracer particles as to self-gravitating populations. As a first example, in Section 3.3 we demonstrated how a subset of dark matter particles chosen to be on radially-biased orbits mix back into the population (Figure 5). This is the case even for halos that, as a whole, are in stable equilibrium; therefore we referred to the phenomenon as a ‘continuous’ radial orbit instability.

The same argument implies that stars undergo the radial orbit instability as easily as dark matter particles. In particular, without a stable disk to enforce extra invariants, dwarf spheroidal galaxies likely have stellar populations with a near-SE distribution (Section 3.4; Figure 6). This provides a footing on which to base spherical Jeans or Schwarzschild analyses of observed systems: it implies an extra prior which can be formulated loosely as stating that $\beta(r) \rightarrow 0$ as $r \rightarrow 0$. Such a prior could be powerful in breaking the degeneracy between density estimates and anisotropy (e.g. Charbonnier et al. 2011). In turn tightening limits on the particle physics of dark matter (Pontzen & Governato 2014).

As a final application, we turned to the question of the baryonic processes that convert a dark matter cusp into a core (Section 3.5). Angular momentum is gained by individual particles during the cusp-flattening process (Tonini et al. 2006) but our earlier work (especially PG12) has focussed primarily on the energy gains instead. In Figure 8 we see that, in spherical or in triaxial dark matter halos, time-dependent perturbations (corresponding to the behaviour of gas in the presence of bursty star formation feedback) always lead to a rise in energy. However, in the spherical case the mean angular momentum as a fraction of j_{circ} drops because the angular momentum of each particle is fixed, whereas j_{circ} rises with E . Only in the triaxial case does $\langle j \rangle / j_{\text{circ}}$ stay constant, indicating a re-isotropisation of the velocities. We tied the increase in $\langle j \rangle$ to the continuous radial orbit instability which always pushes a population with low $\langle j \rangle / j_{\text{circ}}$ back towards the SE limit (Section 3.3).

The consequence is that – in realistically triaxial halos – the final dark matter core size will be chiefly determined by the total energy lost from baryons to dark matter, with little sensitivity to details of the coupling mechanism.

Although the PG12 model can be completed neatly in this way, other analytic calculations or simulations based on spherical symmetry will need to be considered on a case-by-case basis. In our exactly-spherical test cases (Figure 8), core development is substantially suppressed. This certainly cautions against taking the results of purely spherical analyses at face value. On the other hand, any slight asphericity is normally sufficient to prevent the unphysical angular-momentum lock-up – in particular, we verified that the simulations in Teyssier et al. (2013) were unaffected by this issue because their baryons settle into a flattened disk. The analytic result is generic, so the exact shape and strength of the asphericity is a secondary effect in determining the final spherically-averaged distribution function f_0 .

That said, to understand the way in which f_0 actually evolves towards the SE limit (and perhaps freezes out before it gets there) requires going to second order in perturbation theory, as shown in Appendix B. At this point it may also become important to incorporate self-gravity, i.e. the instantaneous connection between δf and δH . This approach has been investigated more fully elsewhere, leading to a different set of insights regarding the onset of the radial orbit instability (e.g. Saha 1991) as opposed to its end state. The present work actually suggests that the radial orbit instability can be cast largely as a kinematical process, and that the self-gravity is a secondary aspect; it would be interesting to further understand how these two views relate. But of more immediate practical importance is to apply the broader insights about dwarf spheroidal stellar equilibria to observational data, something that we will attempt in the near future.

ACKNOWLEDGMENTS

All simulation analysis made use of the pynbody suite (Pontzen et al. 2013). AP acknowledges helpful conversations with James Binney, Simon White, Chervin Laporte and Chiara Tonini and support from the Royal Society and, during 2013 when a significant part of this work was undertaken, the Oxford Martin School. JIR acknowledges support from SNF grant PP00P2_128540/1. FG acknowledges support from HST GO-1125, NSF AST-0908499. NR is supported by STFC and the European Research Council under the European Community’s Seventh Framework Programme (FP7/2007-2013) / ERC grant agreement no 306478-CosmicDawn. JD’s research is supported by Adrian Beecroft and the Oxford Martin School. This research used the DiRAC Facility, jointly funded by STFC and the Large Facilities Capital Fund of BIS.

References

- Adams F. C., Bloch A. M., Butler S. C., Druce J. M., Ketchum J. A., 2007, *ApJ*, 670, 1027, 0708.3101
- Arnold V. I., 1963, *Russian Mathematical Surveys*, 18, 9
- Barnes E. I., Lanzel P. A., Williams L. L. R., 2009, *ApJ*, 704, 372, 0908.3873
- Barnes J., Hut P., Goodman J., 1986, *ApJ*, 300, 112
- Bellovary J. M., Dalcanton J. J., Babul A., Quinn T. R., Maas R. W., Austin C. G., Williams L. L. R., Barnes E. I., 2008, *ApJ*, 685, 739, 0806.3434

- Binney J. J., Tremaine S., 2008, *Galactic dynamics*. 2nd ed. Princeton, NJ, Princeton University Press
- Bonnivard V., Combet C., Maurin D., Walker M. G., 2015, *MNRAS*, 446, 3002, 1407.7822
- Bryan S. E., Mao S., Kay S. T., Schaye J., Dalla Vecchia C., Booth C. M., 2012, *MNRAS*, 422, 1863, 1109.4612
- Charbonnier A., Combet C., Daniel M., Funk S., Hinton J. A., Maurin D., Power C., Read J. I., Sarkar S., Walker M. G., Wilkinson M. I., 2011, *MNRAS*, 418, 1526, 1104.0412
- Cole S., Lacey C., 1996, *MNRAS*, 281, 716, arXiv:astro-ph/9510147
- Corless V. L., King L. J., 2007, *MNRAS*, 380, 149, astro-ph/0611913
- de Zeeuw T., Merritt D., 1983, *ApJ*, 267, 571
- Dehnen W., 1993, *MNRAS*, 265, 250
- Di Cintio A., Brook C. B., Maccio A. V., Stinson G. S., Knebe A., Dutton A. A., Wadsley J., 2014, *MNRAS*, 437, 415, 1306.0898
- Dubinski J., Carlberg R. G., 1991, *ApJ*, 378, 496
- Garrison-Kimmel S., Rocha M., Boylan-Kolchin M., Bullock J., Lally J., 2013, *MNRAS*, 433, 3539, 1301.3137
- Goldstein H., Poole C., Safko J., 2002, *Classical Mechanics*, 3rd edition. Addison Wesley
- Governato F., Zolotov A., Pontzen A., Christensen C., Oh S. H., Brooks A. M., Quinn T., Shen S., Wadsley J., 2012, *MNRAS*, 422, 1231, 1202.0554
- Hayashi E., Navarro J. F., 2006, *MNRAS*, 373, 1117, astro-ph/0608376
- Hayashi E., Navarro J. F., Springel V., 2007, *MNRAS*, 377, 50, arXiv:astro-ph/0612327
- Henon M., 1973, *A&A*, 24, 229
- Hernquist L., 1990, *ApJ*, 356, 359
- Holley-Bockelmann K., Mihos J. C., Sigurdsson S., Hernquist L., 2001, *ApJ*, 549, 862, arXiv:astro-ph/0011504
- Huss A., Jain B., Steinmetz M., 1999, *ApJ*, 517, 64, arXiv:astro-ph/9803117
- Jing Y. P., Suto Y., 2002, *ApJ*, 574, 538, arXiv:astro-ph/0202064
- Kasun S. F., Evrard A. E., 2005, *ApJ*, 629, 781, arXiv:astro-ph/0408056
- Kirk D., 1992, *Graphics Gems III*. AP Professional, Cambridge, MA
- Kolmogorov A. N., 1954, in *Dokl. Akad. Nauk SSSR*, Vol. 98, pp. 527–530
- Lichtenberg A. J., Lieberman M. A., 1992, *Regular and stochastic motion*, 2nd ed
- Loebman S. R., Ivezić Ž., Quinn T. R., Governato F., Brooks A. M., Christensen C. R., Jurić M., 2012, *ApJ*, 758, L23, 1209.2708
- Lynden-Bell D., 1967, *MNRAS*, 136, 101
- MacMillan J. D., Widrow L. M., Henriksen R. N., 2006, *ApJ*, 653, 43, arXiv:astro-ph/0604418
- Maréchal L., Perez J., 2009, *ArXiv e-prints*, 0910.5177
- Martizzi D., Teyssier R., Moore B., Wentz T., 2012, *MNRAS*, 422, 3081, 1112.2752
- Mashchenko S., Couchman H. M. P., Wadsley J., 2006, *Nature*, 442, 539, arXiv:astro-ph/0605672
- Merritt D., Poon M. Y., 2004, *ApJ*, 606, 788, astro-ph/0302296
- Merritt D., Valluri M., 1996, *ApJ*, 471, 82, arXiv:astro-ph/9602079
- Meza A., Zamorano N., 1997, *ApJ*, 490, 136, astro-ph/9707004
- Moser J., 1962, *Nachr. Akad. Wiss. Göttingen Math.-Phys. Kl. II*, 1
- Navarro J. F., Frenk C. S., White S. D. M., 1996, *ApJ*, 462, 563
- , 1997, *ApJ*, 490, 493
- Navarro J. F., Ludlow A., Springel V., Wang J., Vogelsberger M., White S. D. M., Jenkins A., Frenk C. S., Helmi A., 2010, *MNRAS*, 402, 21, 0810.1522
- Oñorbe J., Boylan-Kolchin M., Bullock J. S., Hopkins P. F., Kerš D., Faucher-Giguère C.-A., Quataert E., Murray N., 2015, *MNRAS*, submitted, arXiv:1502.02036
- Peñarrubia J., Pontzen A., Walker M. G., Koposov S. E., 2012, *ApJ*, 759, L42, 1207.2772
- Pontzen A., Governato F., 2012, *MNRAS*, 421, 3464, 1106.0499
- , 2013, *MNRAS*, 430, 121, 1210.1849
- , 2014, *Nature*, 506, 171, 1402.1764
- Pontzen A., Roskar R., Stinson G., Woods R., 2013, *pynbody: N-Body/SPH analysis for python*. Astrophysics Source Code Library ascl:1305.002
- Read J. I., Gilmore G., 2005, *MNRAS*, 356, 107, arXiv:astro-ph/0409565
- Read J. I., Wilkinson M. I., Evans N. W., Gilmore G., Kleya J. T., 2006, *MNRAS*, 367, 387, astro-ph/0511759
- Saha P., 1991, *MNRAS*, 248, 494
- Saha P., Read J. I., 2009, *ApJ*, 690, 154, 0807.4737
- Schneider M. D., Frenk C. S., Cole S., 2012, *J. Cosmology Astropart. Phys.*, 5, 30, 1111.5616
- Stadel J., Potter D., Moore B., Diemand J., Madau P., Zemp M., Kuhlen M., Quilis V., 2009, *MNRAS*, 398, L21, 0808.2981
- Taylor J. E., Navarro J. F., 2001, *ApJ*, 563, 483
- Teyssier R., 2002, *A&A*, 385, 337, arXiv:astro-ph/0111367
- Teyssier R., Pontzen A., Dubois Y., Read J. I., 2013, *MNRAS*, 429, 3068, 1206.4895
- Tonini C., Lapi A., Salucci P., 2006, *ApJ*, 649, 591, arXiv:astro-ph/0603051
- Valluri M., Debattista V. P., Quinn T., Moore B., 2010, *MNRAS*, 403, 525, 0906.4784
- van Albada T. S., 1982, *MNRAS*, 201, 939
- Weinberg M. D., 1991, *ApJ*, 368, 66
- Zolotov A., Brooks A. M., Willman B., Governato F., Pontzen A., Christensen C., Dekel A., Quinn T., Shen S., Wadsley J., 2012, *ApJ*, 761, 71, 1207.0007
- Zorzi A. F., Muzzio J. C., 2012, *MNRAS*, 423, 1955, 1204.5428

APPENDIX A: BACKGROUND

A1 A brief review of actions

This Appendix contains a very brief review of actions which are necessary for deriving the main result in the paper. For more complete introductions see Binney & Tremaine (2008) or Goldstein et al. (2002). We consider any mechanical problem described by phase-space coordinates \mathbf{q} and momenta \mathbf{p} , with Hamiltonian $H(\mathbf{p}, \mathbf{q})$ so that the equations of motion are

$$\dot{\mathbf{p}} = -\frac{\partial H}{\partial \mathbf{q}}, \quad \dot{\mathbf{q}} = \frac{\partial H}{\partial \mathbf{p}}, \quad (\text{A1})$$

where $\dot{\mathbf{q}} = d\mathbf{q}/dt$ and t denotes time.

The actions can be defined starting from any coordinate system in which the motion is separately periodic in each dimension (*i.e.* for each i , q_i and p_i repeat every Δt_i). The actions J_i are then given by

$$J_i \equiv \frac{1}{2\pi} \int_0^{\Delta t_i} p_i \dot{q}_i dt \text{ (no sum over } i\text{)}. \quad (\text{A2})$$




Figure A1. An example orbit in a 3D anisotropic harmonic oscillator. (Left) the solid line shows a portion of the orbit projected into the x - y plane. The equivalent orbit in an isotropic potential is closed and given by the blue dashed line. (Right) the orbit analysed in terms of J_r and j , the actions for the spherical background Hamiltonian. Because ΔH is bounded (see text) the particle has to stay within a narrow band in the J_r - j plane, corresponding to the linear theory resonance (dotted line). The blue dot shows the action for the orbit in the spherical potential (for which the actions are by definition fixed).

By construction the actions J_i do not change under time evolution and are therefore integrals of motion.

We can complete the set of 6 phase-space coordinates with angles Θ in such a way that equations of motion of the form (A1) apply. Since J_i is constant, we must then have

$$0 = \dot{J}_i = -\frac{\partial H}{\partial \Theta_i}, \quad \dot{\Theta}_i = \frac{\partial H}{\partial J_i} \equiv \Omega_i(\mathbf{J}), \quad (\text{A3})$$

where the first equation establishes that H can have no Θ dependence, and the second that consequently the Θ_i each increase at a constant rate in time specified by the frequency $\Omega_i(\mathbf{J})$. The convenience of this set of coordinates is that all the time evolution of a particle trajectory is represented in a very simple way:

$$J_i = \text{constant}, \quad \Theta_i = \text{constant} + \Omega_i t. \quad (\text{A4})$$

Furthermore the equations of motion are canonical, which is sufficient to demonstrate that the coordinates are canonical – in other words, the measure appearing in phase space integrals is $d^3J d^3\Theta$.

We now specialise to the spherical case, with polar coordinates $\mathbf{q} = (r, \theta, \phi)$ and momenta $\mathbf{p} = (\dot{r}, r^2\dot{\theta}, r^2\sin^2\theta\dot{\phi})$. Consider first J_ϕ ,

$$J_\phi = \frac{1}{2\pi} \int_0^{\Delta\phi} dt \dot{\phi}^2 r^2 \sin^2\theta = \frac{1}{2\pi} \int_0^{2\pi} d\phi \dot{\phi} r^2 \sin^2\theta = j_z \quad (\text{A5})$$

where j_z is the z component of the specific angular momentum, $j_z = \dot{\phi} r^2 \sin^2\theta$ which is a constant of motion. Next consider J_θ ,

$$J_\theta = \frac{1}{2\pi} \int_0^{\Delta\theta} dt \dot{\theta}^2 r^2 = \frac{1}{\pi} \int_{\theta_a}^{\theta_b} d\theta \sqrt{j^2 - \frac{j_z^2}{\sin^2\theta}}, \quad (\text{A6})$$

where $j^2 = j_x^2 + j_y^2 + j_z^2$ is the square of the total specific angular momentum, and θ varies between θ_a and θ_b over the course of an orbit. To evaluate the integral requires a change of variables to

relate θ_a and θ_b to the inclination of the orbit and so to j and j_z (e.g. Goldstein et al. 2002, eq 10.135); the final result is that

$$J_\theta = j - j_z. \quad (\text{A7})$$

Because linear combinations of actions are still actions (*i.e.* they still satisfy equation (A3)) one can take j_z as the first and j as the second action in place of J_ϕ and J_θ .

The last action is the radial action,

$$J_r = \frac{1}{2\pi} \int_0^{\Delta r} dt \dot{r}^2 = \frac{1}{\pi} \int_{r_a}^{r_b} dr \sqrt{E - j^2/2r^2 - \Phi(r)}, \quad (\text{A8})$$

where \dot{r} has been evaluated by energy conservation, and r librates between r_a and r_b over the period of an orbit.

Although the complexity of expression (A8) appears to make using the actions cumbersome, the great simplification it brings to the equations of motion in the background (*i.e.* equation A3) makes the perturbation theory tractable. For that reason we have used the action-angle coordinates in our analytic derivation, Section 2, but avoided them when discussing results from simulations in Section 3.

A2 Perturbed trajectories in the harmonic oscillator case

Section 2 used Hamiltonian perturbation theory to discuss the behaviour of particles in aspherical potentials. To connect this more firmly with the equations of motion, it can be helpful to study orbits in a specific potential and connect the solutions with the more general statements made by the perturbation theory. In this Appendix, we use the harmonic oscillator as such an illustrative example.

Consider the Hamiltonian for a single particle in an

anisotropic harmonic oscillator potential,

$$H = \frac{1}{2} (\pi_x^2 + \pi_y^2 + \pi_z^2) + \frac{\omega_0^2}{2} (x^2 + (1 + \varepsilon)y^2 + (1 - \delta)z^2), \quad (\text{A9})$$

where π_x, π_y, π_z denote the momentum in the three cartesian directions and without loss of generality $\varepsilon \geq 0$ and $\delta \geq 0$.

We want to compare the motion in this potential to the behaviour in a sphericalised version with Hamiltonian H_0 ; the change in the Hamiltonian between the true and sphericalised cases is given by

$$\delta H = H - H_0 = \frac{1}{2} \omega_0^2 (\varepsilon y^2 - \delta z^2). \quad (\text{A10})$$

We can immediately read off the first result, which is that the magnitude of δH is bounded. The true solution moves around on a fixed H surface, meaning the fractional error takes a maximum value given by

$$\frac{|\delta H|}{H} < \max(\varepsilon, \delta). \quad (\text{A11})$$

This is the equivalent of the general statement that H_0 variations are small, given by equation (8).

On the other hand, the angular momentum does change significantly. We can see this as follows: each of x, y and z undergoes oscillation at the frequencies $\omega_0, (1 + \varepsilon)^{1/2} \omega_0$ and $(1 - \delta)^{1/2} \omega_0$ respectively. Assuming these new frequencies are not commensurate, the relative phases between the different oscillations slowly shifts until at some point all three separated oscillators reach $x = 0, y = 0$ and $z = 0$ at the same moment. At this point, since the velocity remains finite, the angular momentum has become zero. It may take a number of dynamical times before this happens, but (for example) at the centre of dark matter halos the dynamical time is very short compared to the Hubble time so angular momentum conservation is effectively destroyed.

All the above is illustrated in Figure A1. The left panel contrasts the orbits for a spherical harmonic oscillator (dashed line) and an aspherical oscillator (solid line) projected in the (x, y) plane. The latter obeys equation (A9) with $\varepsilon = \delta = 0.1$ (and the former with $\varepsilon = \delta = 0$). The orbits for the spherical case are closed because the frequencies are identical, so the relative phase of the x and y part of the motion remains fixed. Orbits in the aspherical potential are more complex as the relative phase of the cartesian components gradually changes; in fact, a particle will sometimes plunge through the centre of the potential. This is known as a ‘box orbit’. In more general triaxial potentials, a variety of orbit types are possible (e.g. Merritt & Valluri 1996); the importance of these will be considered momentarily.

The same portion of the orbit is illustrated in the right panel, but now projected into the spherical actions plane (J_r, j) . For the spherical case, J_r and j are exactly conserved by construction and the orbit appears as a single point. In the aspherical case (solid line), J_r and j are not even approximately conserved over more than a dynamical time. However, even then, the orbit remains close to the dashed line of constant H_0 , as required by equation (A11) and more generally by equation (8). From this figure, it is intuitively plausible that the particle is equally likely to be found anywhere along the constant H_0 contour, which is the essential result of Section 2.2.

Another way to look at the effect is to consider the relationship between the angular momentum and the harmonic oscillator’s cartesian actions, J_x, J_y and J_z which remain constant even in the aspherical case. Specialising for simplicity to the case that $J_z = 0$,




Figure A2. Orbits in the action space of the equilibrium ‘Dwarf’ cosmological halo. Dotted contours of constant spherical specific energy H_0 are spaced at $600 \text{ km}^2 \text{ s}^{-2}$. (For this halo $v_{\text{max}}^2 \simeq 3100 \text{ km}^2 \text{ s}^{-2}$.) The straightness of these contours is part of the reason that orbits can efficiently explore the space, as described in the text.

one can show that

$$j = 2 \sqrt{J_x J_y} \sin(\Theta_x - \Theta_y), \quad (\text{A12})$$

where Θ_x and Θ_y are the angles conjugate to J_x and J_y . Only if $\Theta_x - \Theta_y$ remains constant is j an integral of motion; as soon as the oscillator is aspherical, one has

$$j = 2 \sqrt{J_x J_y} \sin(\phi_0 + (\Omega_x - \Omega_y)t) \quad (\text{A13})$$

so that j oscillates on the timescale $2\pi(\Omega_x - \Omega_y)^{-1} \simeq \pi\varepsilon^{-1}$. The situation in 3D is qualitatively similar.

The harmonic oscillator orbits discussed above are all box orbits. More general triaxial potentials support other orbit types (loop orbits, for example, which are much more tightly constrained; or chaotic orbits, which are even less tightly constrained than box orbits). However the Hamiltonian analysis in the main text is general for all these types of possible orbit. The fraction of different orbit types will determine how fast and how far particles diffuse along the contour. In realistic cosmological dark matter halos, most orbits in the central regions are indeed of the box or chaotic type (Zorzi & Muzzio 2012). Even with the baryonic contribution to the gravitational force included (which partially sphericalises the central potential), the large majority of particles remain on the same class of orbit as the dark progenitor (Valluri et al. 2010) so long as feedback is strong enough to prevent long-lived central baryon concentrations developing (Bryan et al. 2012).

A3 The action space of dark matter halos

The action-angle space of our equilibrium ‘Dwarf’ dark-matter-only halo is illustrated in Figure A2, along with some particle orbits (solid lines) over $1.5 \text{ Gyr} \simeq 3t_{\text{dyn}}$. The Hamiltonian is a function of J_r and j only for any spherical potential, and so we have suppressed the third action j_z . As expected, the particles explore the space, approximately running along the H_0 contours, giving rise to the continuous radial orbit instability described in Section 3.3. The contours of H_0 give a great deal of dynamical information, because the background frequencies are defined by $\Omega_0 \equiv \partial H_0 / \partial J$. These

frequencies are at the core of perturbation theory through equation (6), but they also determine the higher-order behaviour as follows.

The first step to understanding the behaviour of resonances is to move to secular perturbation theory (e.g. section 2.4 of Lichtenberg & Lieberman 1992). Secular analysis splits resonant orbits into two classes, known as accidentally- and intrinsically-degenerate. The intrinsic case refers to the situation where the resonance condition applies globally; in other words, that $\mathbf{n} \cdot \boldsymbol{\Omega}_0$ is near-constant along lines of constant H_0 . Suppose, conversely, that $\boldsymbol{\Omega}_0$ did vary along these directions; then, since $\boldsymbol{\Omega}_0$ is defined by the normal to the H_0 contours ($\boldsymbol{\Omega}_0 = \partial H_0 / \partial \mathbf{J}$), this would imply a significant curvature of the dotted lines in Figure A2. Since there is no such curvature, we can read off that the frequencies do not change and the dynamics is in the intrinsically-degenerate regime, giving rise to large-scale migrations. (The same property also means that the approximation $\Omega_r = \Omega_r(H_0)$ used in reaching equation (18) will be extremely accurate.) We have established using numerical investigations that this intrinsic degeneracy property is generic to any spherical action space with smooth potentials, rather than being specific to cosmological halos.

APPENDIX B: THE EVOLUTION OF F_0

In the main text (Section 2), we showed that a distribution function f is maximally stable to linear perturbations if its sphericalised part f_0 appears ergodic, $f_0 = f_0(H_0)$. However we did not discuss the actual evolution of f_0 to see whether this limit is likely to be achieved.

This requires time-dependent perturbation theory. We start by writing the collisionless Boltzmann equation,

$$\frac{\partial f}{\partial t} = [H, f] = [H_0, f_0] + [H_0, \delta f] + [\delta H, f_0] + [\delta H, \delta f], \quad (\text{B1})$$

which is an exact expression. The first term vanishes identically; the second and third terms are linear order, and the final term is second order. The evolution of f_0 is given by taking the time derivative of equation (9) and interchanging the derivative and integral:

$$\begin{aligned} \frac{\partial f_0}{\partial t} &= \frac{1}{(2\pi)^3} \int d^3\Theta [H, f] \\ &= \frac{1}{(2\pi)^3} \sum_{\mathbf{m}, \mathbf{n}} \int d^3\Theta [\delta H_{\mathbf{m}} e^{i\mathbf{m} \cdot \boldsymbol{\Theta}}, \delta f_{\mathbf{n}} e^{i\mathbf{n} \cdot \boldsymbol{\Theta}}], \end{aligned} \quad (\text{B2})$$

where the two linear-order terms have vanished after integrating over $\boldsymbol{\Theta}$. Expanding the Poisson bracket and integrating the remaining term gives

$$\frac{\partial f_0}{\partial t} = -i \sum_{\mathbf{n}} \mathbf{n} \cdot \frac{\partial}{\partial \mathbf{J}} (\delta f_{\mathbf{n}} \delta H_{-\mathbf{n}}). \quad (\text{B3})$$

This shows that the evolution of f_0 is a fundamentally non-linear phenomenon. To make further progress, we can eliminate $\delta f_{\mathbf{n}}$, showing that f_0 evolution depends only on $\delta H_{\mathbf{n}}$. First, Fourier transform the time-dependence of f_0 and δf , so that

$$f_0(\mathbf{J}, t) = \int d\omega e^{i\omega t} \tilde{f}_0(\mathbf{J}, \omega) \quad \text{and} \quad (\text{B4})$$

$$\delta f(\mathbf{J}, \boldsymbol{\Theta}, t) = \sum_{\mathbf{n}} \int d\omega \delta \tilde{f}_{\mathbf{n}}(\mathbf{J}, \omega) e^{i\omega t + i\boldsymbol{\Theta} \cdot \mathbf{n}}. \quad (\text{B5})$$

For simplicity, we will restrict attention to the case where δH is constant in time. Then the evolution of δf is given by the linear-order part of equation (B1), Fourier transformed to give:

$$(\omega + \boldsymbol{\Omega}_0 \cdot \mathbf{n}) \delta \tilde{f}_{\mathbf{n}} = \left(\frac{\partial \tilde{f}_0}{\partial \mathbf{J}} \cdot \mathbf{n} \right) \delta H_{\mathbf{n}}. \quad (\text{B6})$$

This is just the time-dependent, Fourier-transformed version of equation (12). Note that, for δH and δf to be real, the Fourier coefficients must satisfy

$$\delta \tilde{f}_{\mathbf{n}}(\omega)^* = \delta \tilde{f}_{-\mathbf{n}}(-\omega) \text{ and } \delta H_{\mathbf{n}}^* = \delta H_{-\mathbf{n}}. \quad (\text{B7})$$

These requirements are consistent with the relation (B6).

We can now substitute the linear-order solution (B6) for $\delta f_{\mathbf{n}}$ in equation (B3) to get the leading-order evolution equation for f_0 :

$$\omega \tilde{f}_0 = - \sum_{\mathbf{n}} \mathbf{n} \cdot \frac{\partial}{\partial \mathbf{J}} \left[\frac{\delta H_{\mathbf{n}} \delta H_{-\mathbf{n}} \mathbf{n} \cdot \partial \tilde{f}_0 / \partial \mathbf{J}}{\omega + \boldsymbol{\Omega}_0 \cdot \mathbf{n}} \right]. \quad (\text{B8})$$

Using the reality condition, equation (B7), we can pair up negative and positive \mathbf{n} modes, giving an alternative version of the expression that is more explicitly symmetric:

$$\tilde{f}_0 = \sum_{\mathbf{n}} \mathbf{n} \cdot \frac{\partial}{\partial \mathbf{J}} \left[\frac{|\delta H_{\mathbf{n}}|^2 \mathbf{n} \cdot \partial \tilde{f}_0 / \partial \mathbf{J}}{|\boldsymbol{\Omega}_0 \cdot \mathbf{n}|^2 - \omega^2} \right], \quad (\text{B9})$$

where we have divided both sides by ω and so the result is technically only applicable for $\omega \neq 0$. Provided the evolution of f_0 is smooth, its Fourier transform \tilde{f}_0 is continuous, so this is not a problem in practice.

Equation (B9) is enough to give some insight into the relaxation process. For simplicity, consider a perturbation consisting of a single, resonant \mathbf{n}_{\perp} -mode $\delta H_{\mathbf{n}_{\perp}}$ with $\mathbf{n}_{\perp} \cdot \boldsymbol{\Omega}_0 = 0$. Then we can explicitly invert the Fourier transform to yield

$$\frac{\partial^2 f_0}{\partial t^2} = \mathbf{n}_{\perp} \cdot \frac{\partial}{\partial \mathbf{J}} [|\delta H_{\mathbf{n}_{\perp}}|^2 \mathbf{n}_{\perp} \cdot \partial f_0 / \partial \mathbf{J}]. \quad (\text{B10})$$

This is a wave equation in the \mathbf{n}_{\perp} direction with varying speed-of-sound proportional to $|\delta H_{\mathbf{n}_{\perp}}|$. Any unevenness in the resonant directions will flow away at a speed proportional to the strength of the asphericity, showing explicitly that \tilde{f}_0 evolves towards the spherically ergodic limit. We hope to investigate the full behaviour of equation (B9) for a variety of regimes in future work.